

 **PORTAL**
USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

+ "large objects" synchronizing replicating

Feedback Report a problem Satisfaction survey

Terms used **large objects synchronizing replicating**

Found 425 of 177,263

Sort results by Save results to a Binder

Try an [Advanced Search](#)
Try this search in [The ACM Guide](#)

Display results Search Tips
 Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

1 [Eliminating synchronization bottlenecks using adaptive replication](#)

 Martin C. Rinard, Pedro C. Diniz

May 2003 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,
Volume 25 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(826.28 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

This article presents a new technique, adaptive replication, for automatically eliminating synchronization bottlenecks in multithreaded programs that perform atomic operations on objects. Synchronization bottlenecks occur when multiple threads attempt to concurrently update the same object. It is often possible to eliminate synchronization bottlenecks by replicating objects. Each thread can then update its own local replica without synchronization and without interacting with other threads. When ...

Keywords: Atomic operations, commutativity analysis, parallel computing, parallelizing compilers, replication, synchronization

2 [Optimistic replication](#)

 Yasushi Saito, Marc Shapiro

March 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(656.72 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data replication is a key technology in distributed systems that enables higher availability and performance. This article surveys optimistic replication algorithms. They allow replica contents to diverge in the short term to support concurrent work practices and tolerate failures in low-quality communication links. The importance of such techniques is increasing as collaboration through wide-area and mobile networks becomes popular. Optimistic replication deploys algorithms not seen in tradition ...

Keywords: Replication, disconnected operation, distributed systems, large scale systems, optimistic techniques

3

[Eliminating synchronization bottlenecks in object-based programs using adaptive replication](#)

 Martin Rinard, Pedro Diniz
 May 1999 **Proceedings of the 13th international conference on Supercomputing**
 Publisher: ACM Press
 Full text available: [pdf\(1.27 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

4 Munin: distributed shared memory based on type-specific memory coherence 
 J. K. Bennett, J. B. Carter, W. Zwaenepoel
 February 1990 **ACM SIGPLAN Notices , Proceedings of the second ACM SIGPLAN symposium on Principles & practice of parallel programming PPOPP '90**, Volume 25 Issue 3
 Publisher: ACM Press
 Full text available: [pdf\(1.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We are developing Munin, a system that allows programs written for shared memory multiprocessors to be executed efficiently on distributed memory machines. Munin attempts to overcome the architectural limitations of shared memory machines, while maintaining their advantages in terms of ease of programming. Our system is unique in its use of loosely coherent memory, based on the partial order specified by a shared memory parallel program, and in its use of type-specific memory coherence. Ins ...

5 A parallel, real-time garbage collector 
 Perry Cheng, Guy E. Blelloch
 May 2001 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2001 conference on Programming language design and implementation PLDI '01**, Volume 36 Issue 5
 Publisher: ACM Press
 Full text available: [pdf\(1.82 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe a parallel, real-time garbage collector and present experimental results that demonstrate good scalability and good real-time bounds. The collector is designed for shared-memory multiprocessors and is based on an earlier collector algorithm [2], which provided fixed bounds on the time any thread must pause for collection. However, since our earlier algorithm was designed for simple analysis, it had some impractical features. This paper presents the extensions necessary for a pract ...

6 On bounding time and space for multiprocessor garbage collection 
 Guy E. Blelloch, Perry Cheng
 May 1999 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1999 conference on Programming language design and implementation PLDI '99**, Volume 34 Issue 5
 Publisher: ACM Press
 Full text available: [pdf\(1.85 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the first multiprocessor garbage collection algorithm with provable bounds on time and space. The algorithm is a real-time shared-memory copying collector. We prove that the algorithm requires at most $2(R(I + 2/k) + N + 5PD)$ memory locations, where P is the number of processors, R is the maximum reachable space during a computation (number of locations accessible from the root set), N is the maximum number of reachable objects, ...

7 Adaptive software cache management for distributed shared memory architectures 
 John K. Bennett, John B. Carter, Willy Zwaenepoel
 May 1990 **ACM SIGARCH Computer Architecture News , Proceedings of the 17th**

**annual international symposium on Computer Architecture ISCA '90, Volume**

18 Issue 3a

Publisher: ACM PressFull text available: [pdf\(1.10 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An adaptive cache coherence mechanism exploits semantic information about the expected or observed access behavior of particular data objects. We contend that, in distributed shared memory systems, adaptive cache coherence mechanisms will outperform static cache coherence mechanisms. We have examined the sharing and synchronization behavior of a variety of shared memory parallel programs. We have found that the access patterns of a large percentage of shared data objects fa ...

8 A DSA model for data access in self-organizing systems

Dhavy Gantsou

December 2003 **ACM SIGAda Ada Letters , Proceedings of the 2003 annual ACM SIGAda international conference on Ada: the engineering of correct and reliable software for real-time & distributed systems using ada and related technologies SigAda '03, Volume XXIV Issue 1****Publisher:** ACM PressFull text available: [pdf\(186.38 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Data availability is an important issue for self-organizing systems, which include both Peer-to-Peer (P2P) systems and mobile ad hoc networks (MANETs). In P2P systems, the problem of data availability is solved by replicating data across the network. However, this approach wastes resources, and so is not appropriate for MANETs where resource frugality is essential. Unlike P2P systems, MANETs routing protocols require real-time features to cope with a highly dynamic environment, and efficient syn ...

Keywords: distributed synchronization, distributed system annex, mobile data access, real-time data access, self-organizing systems

9 Performance evaluation of the Orca shared-object system

Henri E. Bal, Raoul Bhoedjang, Rutger Hofman, Ceriel Jacobs, Koen Langendoen, Tim Rühl, M. Frans Kaashoek

February 1998 **ACM Transactions on Computer Systems (TOCS), Volume 16 Issue 1****Publisher:** ACM PressFull text available: [pdf\(179.39 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Orca is a portable, object-based distributed shared memory (DSM) system. This article studies and evaluates the design choices made in the Orca system and compares Orca with other DSMs. The article gives a quantitative analysis of Orca's coherence protocol (based on write-updates with function shipping), the totally ordered group communication protocol, the strategy for object placement, and the all-software, user-space architecture. Performance measurements for 10 parallel applications ill ...

Keywords: distributed shared memory, parallel processing, portability

10 Affinity-based management of main memory database clusters

Minwen Ji

November 2002 **ACM Transactions on Internet Technology (TOIT), Volume 2 Issue 4****Publisher:** ACM PressFull text available: [pdf\(553.96 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We study management strategies for main memory database clusters that are interposed between Internet applications and back-end databases as content caches. The task of management is to allocate data across individual cache databases and to route queries to the appropriate databases for execution. The goal is to maximize effective cache capacity and to minimize synchronization cost. We propose an affinity-based management system for main memory database cLusters (*ALBUM*). *ALBUM* executes ea ...

Keywords: Main memory database, clustering, database administration, database cluster, file organization, query affinity, scalability

11 Synchronization mechanisms for SCRAMNet+ systems

 Stephen Menke, Mark Moir, Srikanth Ramamurthy

June 1998 **Proceedings of the seventeenth annual ACM symposium on Principles of distributed computing**

Publisher: ACM Press

Full text available:  pdf(1.35 MB) Additional Information: [full citation](#), [references](#), [index terms](#)



12 Stack allocation and synchronization optimizations for Java using escape analysis

 Jong-Deok Choi, Manish Gupta, Mauricio J. Serrano, Vugranam C. Sreedhar, Samuel P. Midkiff

November 2003 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 25 Issue 6

Publisher: ACM Press

Full text available:  pdf(632.85 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This article presents an *escape analysis* framework for Java to determine (1) if an object is not reachable after its method of creation returns, allowing the object to be allocated on the stack, and (2) if an object is reachable only from a single thread during its lifetime, allowing unnecessary synchronization operations on that object to be removed. We introduce a new program abstraction for escape analysis, the *connection graph*, that is used to establish reachability relationshi ...

Keywords: Connection graphs, escape analysis, points-to graph

13 Replication for web hosting systems

 Swaminathan Sivasubramanian, Michal Szymaniak, Guillaume Pierre, Maarten van Steen September 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 3

Publisher: ACM Press

Full text available:  pdf(374.99 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Replication is a well-known technique to improve the accessibility of Web sites. It generally offers reduced client latencies and increases a site's availability. However, applying replication techniques is not trivial, and various Content Delivery Networks (CDNs) have been created to facilitate replication for digital content providers. The success of these CDNs has triggered further research efforts into developing advanced <i>Web replica hosting systems</i>. These are systems that ...

Keywords: Web replication, content delivery networks

14 Database replication with Slony-I

Ludovic Marcotte



June 2005 **Linux Journal**, Volume 2005 Issue 134

Publisher: Specialized Systems Consultants, Inc.

Full text available: [html\(22.00 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Move up to a highly available cluster without leaving behind the open-source database you trust.

15 Very concurrent mark-&-sweep garbage collection without fine-grain synchronization 

 Lorenz Huelsbergen, Phil Winterbottom

October 1998 **ACM SIGPLAN Notices , Proceedings of the 1st international symposium on Memory management ISMM '98**, Volume 34 Issue 3

Publisher: ACM Press

Full text available: [pdf\(1.36 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe a new incremental algorithm for the concurrent reclamation of a program's allocated, yet unreachable, data. Our algorithm is a variant of mark-&-sweep collection that---unlike prior designs---runs mutator, marker, and sweeper threads concurrently *without* explicit fine-grain synchronization on shared-memory multiprocessors. A global, but infrequent, synchronization coordinates the per-object coloring marks used by the three threads; fine-grain synchronization is achieve ...

16 Real-time replication garbage collection 

 Scott Nettles, James O'Toole

June 1993 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 1993 conference on Programming language design and implementation PLDI '93**, Volume 28 Issue 6

Publisher: ACM Press

Full text available: [pdf\(1.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We have implemented the first copying garbage collector that permits continuous unimpeded mutator access to the original objects during copying. The garbage collector incrementally replicates all accessible objects and uses a mutation log to bring the replicas up-to-date with changes made by the mutator. An experimental implementation demonstrates that the costs of using our algorithm are small and that bounded pause times of 50 milliseconds can be readily achieved.

Keywords: concurrent collection, copying garbage collection, incremental collection, real-time garbage collection, replication

17 Session: Group orientation: a paradigm for modern distributed systems 

 Paulo Veríssimo, Luís Rodrigues

September 1992 **Proceedings of the 5th workshop on ACM SIGOPS European workshop: Models and paradigms for distributed systems structuring**

Publisher: ACM Press

Full text available: [pdf\(650.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Increasing use of distributed systems, with the corresponding decentralization, stimulates the need for structuring activities around groups of participants, for reasons of consistency, user-friendliness, performance and dependability. Although there is a significant number of group communication protocols in the literature, they are penetrating too slowly in operating systems technology. Two important reasons are: the literal interpretation generally made of the end-to-end argument, and the lac ...

18

Techniques for reducing consistency-related communication in distributed shared- 

 **memory systems**

John B. Carter, John K. Bennett, Willy Zwaenepoel

August 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 3

Publisher: ACM Press

Full text available:  pdf(2.86 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Distributed shared memory (DSM) is an abstraction of shared memory on a distributed-memory machine. Hardware DSM systems support this abstraction at the architecture level; software DSM systems support the abstraction within the runtime system. One of the key problems in building an efficient software DSM system is to reduce the amount of communication needed to keep the distributed memories consistent. In this article we present four techniques for doing so: software release consistency; m ...

Keywords: cache consistency protocols, distributed shared memory, memory models, release consistency, virtual shared memory

- 19 **Multidatabase systems: Issues in managing long transactions and large objects in a multidatabase system** 

Gopi Attaluri

November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2**

Publisher: IBM Press

Full text available:  pdf(1.64 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Transactions are a useful and important facility in database systems. Transactions have been incorporated into not only traditional database systems, but also advanced transaction systems (such as file systems (LOCUS), engineering design database systems (CAD/CAM) and general programming languages (Avalon, Argus)). A variety of techniques has been used to support transactions. An environment may be comprised of multiple transaction systems. Writing a global application in such an environment can ...

- 20 **Implementation and performance of Munin** 

 John B. Carter, John K. Bennett, Willy Zwaenepoel

September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles SOSP '91**, Volume 25 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.46 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Munin is a distributed shared memory (DSM) system that allows shared memory parallel programs to be executed efficiently on distributed memory multiprocessors. Munin is unique among existing DSM systems in its use of *multiple consistency protocols* and in its use of *release consistency*. In Munin, shared program variables are annotated with their expected access pattern, and these annotations are then used by the runtime system to choose a consistency protocol best suited to that acc ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

Search Results[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#) [e-mail](#)

Results for "(('large object')<in>metadata) <and> (replicate<in>metadata))<and> (...)"

Your search matched 1 of 1351636 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by Relevance in Descending order.

» [Search Options](#)[View Session History](#)[Modify Search](#)[New Search](#) » [Key](#)

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

Display Format: Citation Citation & Abstract 1. Incremental replication for mobility support in OBIWAN

Veiga, L.; Ferreira, P.;

[Distributed Computing Systems, 2002. Proceedings. 22nd International Confer](#)

2-5 July 2002 Page(s):249 - 256

Digital Object Identifier 10.1109/ICDCS.2002.1022262

[AbstractPlus](#) | Full Text: [PDF\(343 KB\)](#) IEEE CNF[Rights and Permissions](#)[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by